

Amendments to Specification

Please replace paragraph [0096] (the Abstract of the Disclosure) with the following amended paragraph:

[0096] Methods, apparatuses and systems facilitating enhanced classification of network traffic. As discussed above, typical mechanisms that classify network traffic analyze explicitly presented or readily discoverable attributes of individual packets against an application signature, such as a combination of protocol identifiers, port numbers and text strings. The present invention that extends beyond analysis of such explicitly presented packet attributes and holistically analyzes data flows, and in some implementations, related data flows against known application behavior patterns to classify the data flows. Implementations of the present invention facilitate the classification of encrypted or compressed network traffic, or where the higher layer information in the data flows are formatted according to a non-public or proprietary protocol. In one embodiment, the enhanced classification functionality analyzes the behavioral attributes of encrypted data flows against a knowledge base of known application behavior patterns to classify the data flows. In one embodiment, the enhanced classification mechanisms described herein operate seamlessly with other Layer 7 traffic classification mechanisms that operate on attributes of the packets themselves. Implementations of the present invention can be incorporated into a variety of network devices, such as traffic monitoring devices, packet capture devices, firewalls, and bandwidth management devices.

Please replace paragraph [0021] with the following amended paragraph:

[0021] Efficient allocation of network resources, such as available network bandwidth, has become critical as enterprises increase reliance on distributed computing environments and wide area computer networks to accomplish critical tasks. The widely-used Transport Control Protocol (TCP)/Internet Protocol (IP) protocol suite, which implements the world-wide data communications network environment called the Internet and is employed in many local area networks, omits any explicit supervisory function over the rate of data transport over the various devices that comprise the network. While there are certain perceived advantages, this characteristic has the consequence of juxtaposing very high-speed packets and very low-speed packets in potential conflict and produces certain inefficiencies. Certain loading conditions degrade performance of networked applications and can even cause instabilities which could lead to overloads that could stop data transfer temporarily.

Please replace paragraph [0025] with the following amended paragraph:

[0025] A crude form of bandwidth management in TCP/IP networks (that is, policies operable to allocate available bandwidth from a single logical link to network flows) is accomplished by a combination of TCP end systems and routers which queue packets and discard packets when some congestion threshold is exceeded. The discarded and therefore unacknowledged packet serves as a feedback mechanism to the TCP transmitter. Routers support various queuing options to provide for some level of bandwidth management. These options generally provide a rough ability to partition and prioritize separate classes of traffic. However, configuring these queuing options with any precision or without side effects is in fact very difficult, and in some cases, not possible. Seemingly simple things, such as the length of the queue, have a profound

effect on traffic characteristics. Discarding packets as a feedback mechanism to TCP end systems may cause large, uneven delays perceptible to interactive users. Moreover, while routers can slow down inbound network traffic by dropping packets as a feedback mechanism to a TCP transmitter, this method often results in retransmission of data packets, wasting network traffic and, especially, inbound capacity of a Wide Area Network (WAN) link. In addition, routers can only explicitly control outbound traffic and cannot prevent inbound traffic from over-utilizing a WAN link. A 5% load or less on outbound traffic can correspond to a 100% load on inbound traffic, due to the typical imbalance between an outbound stream of acknowledgments and an inbound stream of data.

Please replace paragraph [0026] with the following amended paragraph:

[0026] In response, certain data flow rate control mechanisms have been developed to provide a means to control and optimize efficiency of data transfer as well as allocate available bandwidth among a variety of business enterprise functionalities. For example, U.S. 6,038,216 discloses a method for explicit data rate control in a packet-based network environment without data rate supervision. Data rate control directly moderates the rate of data transmission from a sending host, resulting in just-in-time data transmission to control inbound traffic and reduce the inefficiencies associated with dropped packets. Bandwidth management devices allow for explicit data rate control for flows associated with a particular traffic classification. For example, U.S. 6,412,000, above, discloses automatic classification of network traffic for use in connection with bandwidth allocation mechanisms. U.S. Pat. No. 6,046,980 discloses systems and methods allowing for application layer control of bandwidth utilization in

packet-based computer networks. For example, bandwidth management devices allow network administrators to specify policies operative to control and/or prioritize the bandwidth allocated to individual data flows according to traffic classifications. In addition, certain bandwidth management devices, as well as certain routers, allow network administrators to specify aggregate bandwidth utilization controls to divide available bandwidth into partitions. With some network devices, these partitions can be configured to ensure a minimum bandwidth and/or cap bandwidth as to a particular class of traffic. An administrator specifies a traffic class (such as File Transfer Protocol (FTP) data, or data flows involving a specific user) and the size of the reserved virtual link—i.e., minimum guaranteed bandwidth and/or maximum bandwidth. Such partitions can be applied on a per-application basis (protecting and/or capping bandwidth for all traffic associated with an application) or a per-user basis (controlling, prioritizing, protecting and/or capping bandwidth for a particular user). In addition, certain bandwidth management devices allow administrators to define a partition hierarchy by configuring one or more partitions dividing the access link and further dividing the parent partitions into one or more child partitions. While the systems and methods discussed above that allow for traffic classification and application of bandwidth utilization controls on a per-traffic-classification basis operate effectively for their intended purposes, they possess certain limitations. As discussed more fully below, identification of traffic types associated with data flows traversing an access link involves the application of matching criteria or rules to explicitly presented or readily discoverable attributes of individual packets against an application signature which may comprise a protocol identifier (e.g., TCP, HyperText Transport Protocol (HTTP), User Datagram Protocol (UDP), Multipurpose Internet Mail Extensions (MIME) types,

etc.), a port number, and even an application-specific string of text in the payload of a packet. After identification of a traffic type corresponding to a data flow, a bandwidth management device associates and subsequently applies bandwidth utilization controls (e.g., a policy or partition) to the data flow corresponding to the identified traffic classification or type. Accordingly, simple changes to an application, such as a string of text appearing in the payload or the use of encryption text may allow the application to evade proper classification and corresponding bandwidth utilization controls or admission policies.